

[illegible]

Title: METHOD AND APPARATUS FOR ALIGNING MEMORY WRITE DATA IN A MICROPROCESSOR

Filed: August 28, 2001

F. CHAU & ASSOCIATES, LLP  
1900 Hempstead Turnpike, Suite 501  
East Meadow, New York 11554  
Tel.: (516) 357-0091  
Fax : (516) 357-0092

METHOD AND APPARATUS FOR ALIGNING MEMORY WRITE DATA  
IN A MICROPROCESSOR

BACKGROUND OF THE INVENTION

5

1. Technical Field

The present invention relates generally to microprocessors and, in particular, to a method and apparatus for aligning memory write data in a microprocessor.

10

2. Description of Related Art

Contemporary microprocessors usually support the processing of data of multiple bit widths. In such microprocessors, one or more data types are preferred in that they are supported by appropriately-sized hardware primitives, such as registers, arithmetic logic units (ALUs), and memory queues.

15

Other data types are supported to a limited degree, in that such data can be read from and written to their natural alignment boundaries. During a read process, this usually requires realigning into processor registers and processor data paths such that a single data item which is not of full width is naturally aligned at the least significant

20

position, optionally with zero or sign extension. Likewise, a sub-width data type when written to an address of its natural alignment boundary needs to be realigned within a processor to that boundary.

5           Also, it should be noted that data is organized in modern computers in one of two formats or in some combination of two formats. The formats are known as "big-endian", in which the high order bit, byte or other units of information is located in the lower numbered unit address, and "little-endian", in which the high order bit, byte or other units of information is located in the higher numbered unit address.

10           To date, microprocessors have usually included a rotator/alignment network prior to the memory interface for writing data, and beyond the memory stage for the read operation. This is undesirable for at least the following reasons. One such reason concerns data flow considerations. That is, if the standard rotator is to be used, then the data flow from the rotator to the memory interface is complex and can be slow. In particular, the rotator is aligned at ALU track pitch which differs from memory track pitch. Therefor, it is often necessary to introduce a second and possibly third rotator, before (for writing) and

after (for reading) the memory access. This requires additional hardware, and also increases latency of memory operations, including those which do not need alignment operations (usually data in one of the preferred data formats corresponding to a natural machine processing width), since bypassing the rotators is often not practical.

Attempts have been made to reduce the complexity and resultant latency of memory access in the MIPS-X processor prototype from Stanford University, and the initial Alpha processor specification by Digital Equipment Corporation.

The Stanford MIPS processor is described by J. Hennessy, in "VLSI Processor Architecture", IEEE Transactions on Computers, Vol. C-33, No. 12, pp. 1221-46, Dec. 1984. This processor uses byte insert (IC) and extract (XC) operations to manipulate bytes, but otherwise only supports word addressing.

The approach used in these processors was to support only preferred data width memory operations, in conjunction with explicit memory alignment operations. Both of these architecture specifications have only had limited success with this approach, prompting the addition of sub-word memory operations to later generations of the processors.

The usual alignment networks are endian-specific and adding endian-independence usually requires additional logic in the alignment network. Processor implementors can either decide to support both big-endian and little-endian modes at the cost of high complexity, or only support one mode at the cost of sacrificing compatibility with a significant number of processors not having the selected endianness.

Explicit software-based alignment does not suffer from this defect, as both little-endian and big-endian configurations can be supported by the appropriate software sequences.

Accordingly, it would be desirable and highly advantageous to have a method and apparatus that supports software-based alignment of memory accesses, so as to reduce microprocessor implementation complexity, support big-endian and little-endian configurations, and reduce the penalty for using software based alignment of memory-write data found in previous processors.

#### SUMMARY OF THE INVENTION

The problems stated above, as well as other related problems of the prior art, are solved by the present

invention, a method and apparatus for aligning memory write data in a microprocessor.

The present invention advantageously reduces microprocessor implementation complexity, supports  
5 big-endian and little-endian configurations, and reduces the penalty for using software based alignment of memory-write data found in previous processors.

According to an aspect of the present invention, there is provided a method for aligning and inserting data  
10 elements into a memory based upon an instruction sequence consisting of one or more alignment instructions and a single store instruction. Given a data item that includes a data element to be stored, the method includes the step of aligning the data element in another memory with respect to  
15 a predetermined position in the memory, in response to the one or more alignment instructions. A mask is dynamically generated to enable writing of memory bit lines that correspond to the aligned data element. The memory bit lines are written to the memory under a control of the mask.  
20 The generating and writing steps are performed in response to the single store instruction.

According to another aspect of the present invention, there is provided a system for aligning and inserting data

elements into a memory in response to an instruction  
sequence consisting of one or more alignment instructions  
and a single store instruction. The system includes means  
for receiving a data item that includes a data element to be  
5 stored. Further, the system includes means for aligning the  
data element in another memory with respect to a  
predetermined position in the memory, in response to the one  
or more alignment instructions. Also, the system includes  
means for dynamically generating a mask to enable writing of  
10 memory bit lines that correspond to the aligned data  
element, in response to the single store instruction.  
Moreover, the system includes means for writing the memory  
bit lines to the memory under a control of the mask, in  
response to the single store instruction.

15 According to yet another aspect of the present  
invention, there is provided a method for storing data in a  
memory based upon an instruction sequence consisting of one  
or more alignment instructions and a single store  
instruction. The method includes the step of aligning the  
20 data in a register relative to a location of the data within  
a target memory address line, in response to the one or more  
alignment instructions. A portion of the aligned data is  
stored within the memory under a control of data type

information and an address argument specified by the single store instruction, in response to the single store instruction.

According to still yet another aspect of the present invention, there is provided, in a memory system with a register and a memory and without an alignment network, a method for storing data in the memory based upon a single store instruction. The method includes the step of generating memory write instructions for directly storing a plurality of bits from the register to a memory line in the memory under a control of a mask, a bit range of the plurality of bits being specified by parameters of the single store instruction.

These and other aspects, features and advantages of the present invention will become apparent from the following detailed description of preferred embodiments, which is to be read in connection with the accompanying drawings.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

FIG. 1 is a block diagram illustrating a memory write data path used in a processor supporting byte, half-word, and word memory write operations, according to the prior art;



FIG. 2 is a block diagram illustrating the memory write data path used in a processor without support for subword alignment, according to the prior art;

FIG. 3 is a flow diagram illustrating the control logic used in conjunction with the memory write path of FIG. 1, according to prior art;

FIG. 4 is a flow diagram illustrating the control logic used in conjunction with FIG. 2 in the Stanford MIPS and DEC Alpha processors, according to the prior art;

FIG. 5 is a flow diagram illustrating alignment network control logic used in conjunction with the memory write data path of FIG. 2, according to an illustrative embodiment of the present invention;

FIG. 6 is a block diagram illustrating an alignment network, according to an illustrative embodiment of the present invention;

FIG. 7 is a diagram illustrating a 3-instruction store sequence, according to an illustrative embodiment of the present invention;

FIG. 8 is a flow diagram illustrating steps corresponding to a rotate bytes left (rotbl) instruction of FIG. 7, according to the prior art;

FIG. 9 is a flow diagram illustrating steps corresponding to a store byte (stb) instruction of FIG. 7, according to an illustrative embodiment of the present invention;

5        FIG. 10 is a diagram illustrating an instruction sequence for storing 32 bit words from a wide vector register embodiment as shown in FIG. 6 in conjunction with control logic as shown in FIG. 5, according to an illustrative embodiment of the present invention;

10        FIG. 11 is a flow diagram illustrating steps corresponding to a rotate bytes left (vrotbl) instruction of FIG. 10, according to the prior art; and

15        FIG. 12 is a flow diagram illustrating steps corresponding to a store word (vstw) instruction of FIG. 10, according to an illustrative embodiment of the present invention.

#### DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

20        It is to be understood that the present invention may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. Preferably, the present invention is implemented as a combination of both hardware and software, the software

being an application program tangibly embodied on a program storage device. The application program may be uploaded to, and executed by, a machine comprising any suitable architecture. Preferably, the machine is implemented on a computer platform having hardware such as one or more central processing units (CPU), a random access memory (RAM), and input/output (I/O) interface(s). The computer platform also includes an operating system and microinstruction code. The various processes and functions described herein may either be part of the microinstruction code or part of the application program (or a combination thereof) which is executed via the operating system. In addition, various other peripheral devices may be connected to the computer platform such as an additional data storage device.

It is to be further understood that, because some of the constituent system components depicted in the accompanying Figures may be implemented in software, the actual connections between the system components may differ depending upon the manner in which the present invention is programmed. Given the teachings herein, one of ordinary skill in the related art will be able to contemplate these

and similar implementations or configurations of the present invention.

FIG. 1 is a block diagram illustrating the memory write data path 100 used in a processor supporting byte, half-word, and word memory write operations, according to the prior art. The memory write path 100 includes a register file 110, multiplexers 120, and a cache memory 130. It is to be appreciated that while a cache is shown in the FIGs. herein, main memory (not shown) may be used in place of the cache; furthermore, such main memory may be augmented by the cache.

Processors supporting byte, half-word and word memory write operations employ an alignment network which steers the byte, half-word, word, or other unit, to its appropriate location within a memory line. In FIG. 1, the memory line has been designated as a cache line, since most computer systems employ wide caches. However, as is readily apparent to one of ordinary skill in the related art, other memory structures could interface to the alignment network. Moreover, as is known to those of ordinary skill in the related art, a computer system can also include other logic, such as a read/write buffer, Error Correcting Code (ECC) or parity checking, and generation logic. As is evident from

FIG. 1, using such an alignment network requires the use of multiplexers 120 to select the data source to be written to the cache memory 130 depending on control signals generated based on the memory address and data type, resulting in increased chip size and signal delay. Furthermore, the multitude of signals supplied to the multiplexers 120 requires chip area to connect the signals which also increases signal delay.

FIG. 2 is a block diagram illustrating the memory write data path 200 used in a processor without support for subword alignment, according to the prior art. The memory write data path 200 includes a register file 210, multiplexers 220, and a cache memory 230. As can be seen, each byte of the register file 210 is connected only to a memory location corresponding to a natural alignment. Thus, the first byte of the register file 210 can be supplied to a first, fifth, ninth, thirteenth byte of a memory line, the second byte to the second, sixth, tenth, fourteenth byte of a memory line, and so forth. It should be noted that this leads to a simpler connection network and to reduced width multiplexers selecting only between the previous memory content and the newly supplied content. Furthermore, the functionality of the multiplexers 220 can also be achieved

with the use of write-enable control signals to memory banks, thereby completely eliminating the needs for multiplexers and further reducing the complexity of the alignment network.

5 In FIG. 2, the memory line has been designated as a cache line, since most computer systems employ wide caches. However, as is readily apparent to one of ordinary skill in the related art, other memory structures could interface to the alignment network. Moreover, as is known to those of  
10 ordinary skill in the related art, a computer system can also include other logic, such as a read/write buffer, Error Correcting Code (ECC) or parity checking, and generation logic.

FIG. 3 is a flow diagram illustrating the alignment  
15 network control logic 300 used in conjunction with the memory write path 100 of FIG. 1, according to prior art. In particular, this control logic computes output SELECT[BYTE] which selects how the multiplexers 120 select the output data from their input data for the respective bytes. As is  
20 apparent to those of ordinary skill in the related art, other control logic may also be present in the memory data path, e.g., cache hit/miss logic. The alignment network control logic 300 is expressed as a sequential loop to

express the functionality of the alignment network control logic 300, although it is apparent to those of ordinary skill in the related art that actual implementations are usually performed using parallel control logic. In particular, when a store instruction is processed by the processor, the alignment network control logic 300 performs a loop for every byte of the memory line size (MLS), beginning with byte 0 (initialized in step 310) until the end of the memory line size (loop counter increment is performed in step 350 and loop iteration test in step 360).

In step 320, it is determined whether the current byte within a memory line (a) corresponds to the alignment specified by the store address and (b) falls within the data size range specified by the data width. If not (negative determination at step 320), then processing proceeds to step 340. Otherwise, processing continues with step 330.

In step 330, the control signal for the current byte is computed as the control signal corresponding to the byte of the register file to be written to the current location according to the address and data width specified, and processing proceeds to step 350.

In step 340, the multiplexer is configured to maintain the previous value in the memory location, and processing proceeds to step 350.

5 The loop counter is incremented in step 350, and it is determined whether the current byte is less than the MLS (step 360). If so, then processing returns to step 320 to perform the next iteration. The loop iterates until control signals for the entire memory line have been generated. Thus, for a negative determination at step 360, the  
10 processing is terminated.

FIG. 4 is a flow diagram illustrating the alignment network control logic 400 used in conjunction with the memory write data path 200 of FIG. 2 in the Stanford MIPS and, in principle, DEC Alpha processors, according to the  
15 prior art. In particular, the alignment network control logic 400 is used to control multiplexers 220. As is apparent to those of ordinary skill in the related art, other control logic may also be present in the memory data path, e.g., cache hit/miss logic. The alignment network  
20 control logic 400 is expressed as a sequential loop to express the functionality of the alignment network control logic 400, although it is apparent to those of ordinary skill in the related art that actual implementations are



usually performed using parallel control logic. In particular, when a store instruction is processed by the processor, the alignment network control logic 400 performs a loop for every byte of the memory line size (MLS),  
5 beginning with byte 0 (initialized in step 410) until the end of the memory line size (loop counter increment is performed in step 450 and loop iteration test in step 460).

In step 420, it is determined whether the current byte within a memory line is part of the memory word to be  
10 written at the specified store address. Since the Stanford MIPS only supports word write operations, the alignment network control logic 400 is significantly simplified. In particular, the alignment network control logic 400 can be implemented by comparing a subset of the address lines with  
15 a predefined constant. If the determination performed at step 420 returns a negative response, then processing proceeds to step 440. Otherwise, processing continues with step 430.

In step 430, the control signal for the current byte is  
20 selected such that the value provided by the memory write data path will be selected, and processing proceeds to step 450.

In step 440, the multiplexer is configured to maintain the previous value in the memory location, and processing proceeds to step 450.

The loop counter is incremented in step 450, and it is determined whether the current byte is less than the MLS (step 460). If so, then processing returns to step 420 to perform the next iteration. The loop iterates until control signals for the entire memory line have been generated. Thus, for a negative determination at step 460, the processing is terminated.

FIG. 5 is a flow diagram illustrating the alignment network control logic 500 used in conjunction with the memory write data path 200 of FIG. 2, according to an illustrative embodiment of the present invention. In particular, the alignment network control logic 500 is used to control multiplexers 220. As is apparent to those of ordinary skill in the related art, other control logic may also be present in the memory data path, e.g., cache hit/miss logic. The alignment network control logic 500 is expressed as a sequential loop to express the functionality of the alignment network control logic 500, although it is apparent to those of ordinary skill in the related art that actual implementations are usually performed using parallel

control logic. In particular, when a store instruction is processed by the processor, the alignment network control logic 500 performs a loop for every byte of the memory line size (MLS), beginning with byte 0 (initialized in step 510) until the end of the memory line size (loop counter increment is performed in step 550 and loop iteration test in step 560).

In step 520, it is determined whether the current byte within a memory line is part of the memory word to be written at the specified store address. If not (negative determination at step 520), then processing proceeds to step 540. Otherwise, processing continues with step 530.

In step 530, the control signal for the current byte is selected such that the value provided by the memory write data path will be selected, and processing proceeds to step 550.

In step 540, the multiplexer is configured to maintain the previous value in the memory location, and processing proceeds to step 550.

The loop counter is incremented in step 550, and it is determined whether the current byte is less than the MLS (step 560). If so, then processing returns to step 520 to perform the next iteration. The loop iterates until control

signals for the entire memory line have been generated.  
Thus, for a negative determination at step 560, the  
processing is terminated.

FIG. 6 is a block diagram illustrating an alignment  
network 600, according to an illustrative embodiment of the  
present invention. The alignment network 600 is designed to  
further reduce design complexity. In particular, the  
alignment network 600 includes a wide register file 610  
(e.g., corresponding to the width of the memory line size)  
connected to the memory write data path. The alignment  
network further includes multiplexers 620 and a cache memory  
630. The alignment network 600 is to be used in conjunction  
with the alignment network control logic 500 of FIG. 5. By  
employing a wide register file, this embodiment simplifies  
the laying out of signals during the design process, and  
furthermore reduces fan-out of the signals directed at the  
memory write path implementation.

In FIG. 6, the memory line has been designated as a  
cache line, since most computer systems employ wide caches.  
However, as is readily apparent to one of ordinary skill in  
the related art, other memory structures could interface to  
the alignment network. Moreover, as is known to those of  
ordinary skill in the related art, a computer system can

also include other logic, such as a read/write buffer, Error Correcting Code (ECC) or parity checking, and generation logic.

Comparing the instruction sequences to access memory in sub-word data widths, most current microprocessors are based on an alignment network as illustrated in FIG. 1 with control logic as illustrated in FIG. 3. Using such hardware mechanisms, a byte can, for example, be written to main memory with a single instruction such as the PowerPC instruction:

```
stb r4, 0(r5)
```

The preceding instruction is used to store the least significant byte contained in register r4 at the location specified by register r5. It should be noted that the execution of the store byte stb memory instruction implemented in common microprocessor architectures is complex and induces significant design complexity and cycle time penalty. In particular, all memory instructions need to be routed through an alignment network, to be properly aligned relative to a memory line, and then written into the memory line. Furthermore, the store address is checked for

alignment on natural data item boundary during the execution  
of the stb instruction and, if misalignment is detected,  
then complex control logic to store unaligned data items is  
invoked. Storing unaligned data items may usually be  
5 accomplished by transferring control to either a state  
machine or to microcode contained in the processor.

Furthermore, this complexity and cycle time penalty  
impacts all memory store instructions, thus even when  
storing properly aligned data items (such as full width data  
10 items), the alignment network needs to be used. In addition,  
even store instructions storing to addresses known to be  
aligned on natural data item boundaries are usually checked  
for such alignment, incurring additional overhead penalties.

As in the previous example, the significant byte of the  
15 value to be stored is contained in the least significant  
byte of register r4 and the target address specified by  
0(r5). Referring now to an architecture which does not  
support memory access at a smaller than word granularity  
(such as the DEC Alpha Architecture), in the prior art,  
20 storing a byte consists of sequences of loading the  
surrounding memory word, merging the byte with the memory  
word and storing back the merged word, e.g., using the  
following code:

5	<pre> LDQ_U r3, 0(r5)                ;; load preferred length data                                 item (i.e., the full width                                 item) surrounding the                                 sub-width item </pre>
10	<pre> INSB r4,r5,r4                  ;; rotate the sub-width item                                 into the appropriate slot of                                 the full width item </pre>
15	<pre> MSBL r3=r3&amp;~(255&lt;&lt;r5)          ;; clear the old data in the                                 desired sub-width slot of the                                 preferred length data item </pre>
20	<pre> OR    r3, r3, r4               ;; merge the correctly aligned                                 sub-width item into the                                 previously cleared slot </pre>
25	<pre> STQ_U r3, 0(r5)                ;; store back full-width item </pre>

We observe that significant effort is expended in loading, clearing and inserting data from/into a preferred data width item. Similar effort is expended when other variations and similar implementations of the preceding prior art instruction sequence are employed. The invention is equally applicable in overcoming the deficiencies of such other variations and implementations.

An implementation of the present invention that uses the alignment network of FIG. 2 in conjunction with the alignment network control logic 400 of FIG. 4 will now be described. According to the present invention, only the following steps are necessary: rotate the data of the sub-width item into the appropriate slot of the full width item; and store the data into the memory, using byte-wise write-enable functionality in the memory interface (e.g., a read-write buffer).

In one implementation according to the present invention, this sequence can be generated using two specialized instructions. In another implementation according to the present invention, the data rotation can be performed using general purpose instructions, for example, consisting of one instruction to compute the byte (or halfword, or word, etc.) slot within a full-data width register based on the alignment of the specified byte (or halfword, or word, etc.) address, and a second rotation instruction.

Consider the 3-instruction store sequence depicted in FIG. 7. FIG. 7 is a diagram illustrating a 3-instruction store sequence, according to an illustrative embodiment of the present invention.



The functionality of the rotbl (rotate bytes left) instruction of FIG. 7 is depicted in FIG. 8. FIG. 8 is a flow diagram illustrating steps corresponding to the rotate bytes left (rotbl) instruction of FIG. 7, according to the prior art.

Register r3 is rotated to the left by the byte count specified by the 2 rightmost bits of register r3 (that is, by  $(r3 \& 3) * 8$  bits) (step 810). The result is stored in register r2 (step 820).

The functionality of the stb (store byte) instruction of FIG. 7 is depicted in FIG. 9. FIG. 9 is a flow diagram illustrating steps corresponding to the store byte (stb) instruction of FIG. 7, according to an illustrative embodiment of the present invention.

A write mask is computed from the address argument supplied (the sum of displacement 0 and register r5) according to the control logic of FIG. 5 (i.e., the address corresponding to a store instruction initially provided to store the data item) (step 910). For the purposes of the present invention, the terms "mask" and "write mask" refer to a control word that selects data bits, bytes, or words (or some other data type, as appropriate) from a collection of first and second (and optionally, third, fourth, etc.) data collections of bits, bytes, or words, selecting for each resultant bit, byte, or word (or other data type, as appropriate), a bit, byte, or word from either the respective value in the first or the second (and optionally, a third, fourth, etc.) data collections under control of the control word. The mask is formed based on the store datatype to facilitate the writing of a number of bits corresponding to the size of the data type. Thus, if the mask encodes bits, then the number of bits set in the mask will be set to the number bits in the data type in a common encoding of masks. In another encoding, a complemented mask may be used based on the ones-complement of the write mask. It is to be appreciated that the present invention is not limited to the preceding encodings and, thus, other

encodings may be employed in accordance with the present invention while maintaining the spirit and scope thereof.

The word in r2, which contains a correctly aligned data item in a preferred-data-width register r2, is supplied to the memory interface for storage under the control of the computed write mask (step 920).

By combining instructions 1 and 2 into a single rotate instruction which computes the alignment similar to the INSB instruction, the same operation can be achieved with only two instructions.

The description and Figures have been presented based on atomic byte-granularity memory units. However, it will be readily apparent to one of ordinary skill in the related art that alignment and write signal selection can be performed at any other granularity level, both coarser and finer, such as bits, half-bytes, words, half-words, double-words, or any other unit.

Further, it will be readily apparent to one of ordinary skill in the related art that although the present invention has been described in terms of a big-endian memory layout, the present invention can be readily adapted for use in little-endian systems or in systems supporting big- and little-endian processing.

Moreover, it will be readily apparent to one of ordinary skill in the related art that the present invention is also applicable to a processor containing wide vector registers when processing 32-bit words. In particular, 32 bit words can be stored from a wide vector register embodiment as shown in FIG. 6 in conjunction with control logic as shown in FIG. 5 using an instruction sequence similar to that depicted in FIG. 10. The particular code sequence shown therein corresponds to storing a 32-bit word right-aligned within a wide vector register, e.g., in an architecture supporting a preferred slot for processing scalar data in its rightmost vector element. The "preferred slot" is described in a U.S. Patent application, Attorney Docket YOR9-2001-0675, entitled "Processor Implementation Having Unified Scalar and SIMD Datapass, filed on August 14, 2001, the disclosure of which is commonly assigned and incorporated by reference herein.

The functionality of the vrotbl (rotate bytes left) operation of FIG. 10 is depicted in FIG. 11. FIG. 11 is a flow diagram illustrating steps corresponding to the rotate bytes left (vrotbl) instruction of FIG. 10, according to the prior art. Register r3 is rotated to the left by the byte count specified by the 4 rightmost bits of register r3 (that

is, by  $(r3 \& 15) * 8$  bits) (step 1110). The result is stored in register r2 (step 1120).

The functionality of the vstw (store word) operation of FIG. 10 is depicted in FIG. 12. FIG. 12 is a flow diagram illustrating steps corresponding to the store word (vstw) instruction of FIG. 10, according to an illustrative embodiment of the present invention.

A write mask is computed from the address argument supplied (the sum of displacement 0 and register r5) according to the control logic of FIG. 5 (i.e., the address corresponding to a store instruction initially provided to store the data item) (step 1210). The word in r2, which contains a correctly aligned word in a preferred-data-width register r2, is supplied to the memory interface for storage under the control of the computed write mask (step 1220).

It is to be appreciated that the present invention can be used in conjunction with a number of implementation techniques for implementing systems which can access both little and big endian data. Such techniques can be exploited by adding hardware capabilities or (preferably) appropriate software instructions to the code sequence to perform techniques familiar to those of ordinary skill in

related art (e.g., "bit swizzling") in conjunction with the teachings of the present invention.

A sample implementation sequence of bit swizzling used to perform a little-endian store byte operation (for a word-oriented architecture) is as follows, based on Fig 7 in conjunction with a first bit swizzling step. As in previous examples, the value to stored is contained within the least significant byte of register r4, and the memory store address is 0(r5):

```
10      xori r1, r5, 3          ; swizzle r5 into r1
      subfi r3, r1, 3          ; compute alignment
                                   position
      rotbl r2, r4, r3         ; alignment operation
15      stb r2, 0(r1)          ; perform store using mask
```

The apparatus and methods disclosed herein can be used in conjunction with a number of implementation techniques for implementing systems which can access unaligned data.

In particular, this can be achieved with specialized store instructions which perform a store operation under the control of a mask which selects the bits in a first or a second half of an unaligned word, the first half consisting of those data elements to be stored at an address below an

alignment boundary, the second half to be stored at an address above an alignment boundary.

A sample implementation sequence of unaligned store access is provided here based on FIG 10:

```
5      subfi r3, r5, 12          ; compute rotate count
      vrotbl r2, r4, r3         ; rotate to alignment
      vstwl r2, 0(r5)           ; store lower half of word
                                   starting at 0(r5) up to
10      next natural alignment
                                   boundary
      vstwu r2, 0(r5) ;         ; store upper half of word
                                   starting at alignment
                                   boundary to end of word
```

15 When an aligned data item is stored by the instruction sequence directed towards support for storing unaligned data items, at least one of the vstwl and vstwu can be a no-op. As is known, a no-op is an instruction which does nothing.

20 The alignment boundary can be word size, wide word size, memory line size, cache line size, or a similar natural architectural boundary. In this exemplary implementation, the vstwu instruction adjusts the address base to address the second memory portion automatically. In another

implementation, this may have to be done explicitly, e.g.,  
using an instruction such as vstwu r2, 16(r5) in the case of  
a 16-byte alignment boundary. Given the teachings of the  
present invention provided herein, one of ordinary skill in  
5 the related art will contemplate these and various other  
implementations and variations of the present invention,  
while maintaining the spirit and scope of the present  
invention.

Although the illustrative embodiments have been  
10 described herein with reference to the accompanying  
drawings, it is to be understood that the present system and  
method is not limited to those precise embodiments, and that  
various other changes and modifications may be affected  
therein by one skilled in the art without departing from the  
15 scope or spirit of the invention. All such changes and  
modifications are intended to be included within the scope  
of the invention as defined by the appended claims.